

I couldn't find
these documents in the
bibliography. All those
in this box have been
checked.

-A

Tom: CITE THESE
AS ~~THE~~ NOTES OF
The International
Network W'kg.
Group.

ARPA COMPUTER NETWORK PROVIDES COMMUNICATIONS TECHNOLOGY
FOR COMPUTER/COMPUTER INTERACTION WITHIN SPECIAL RESEARCH COMMUNITY

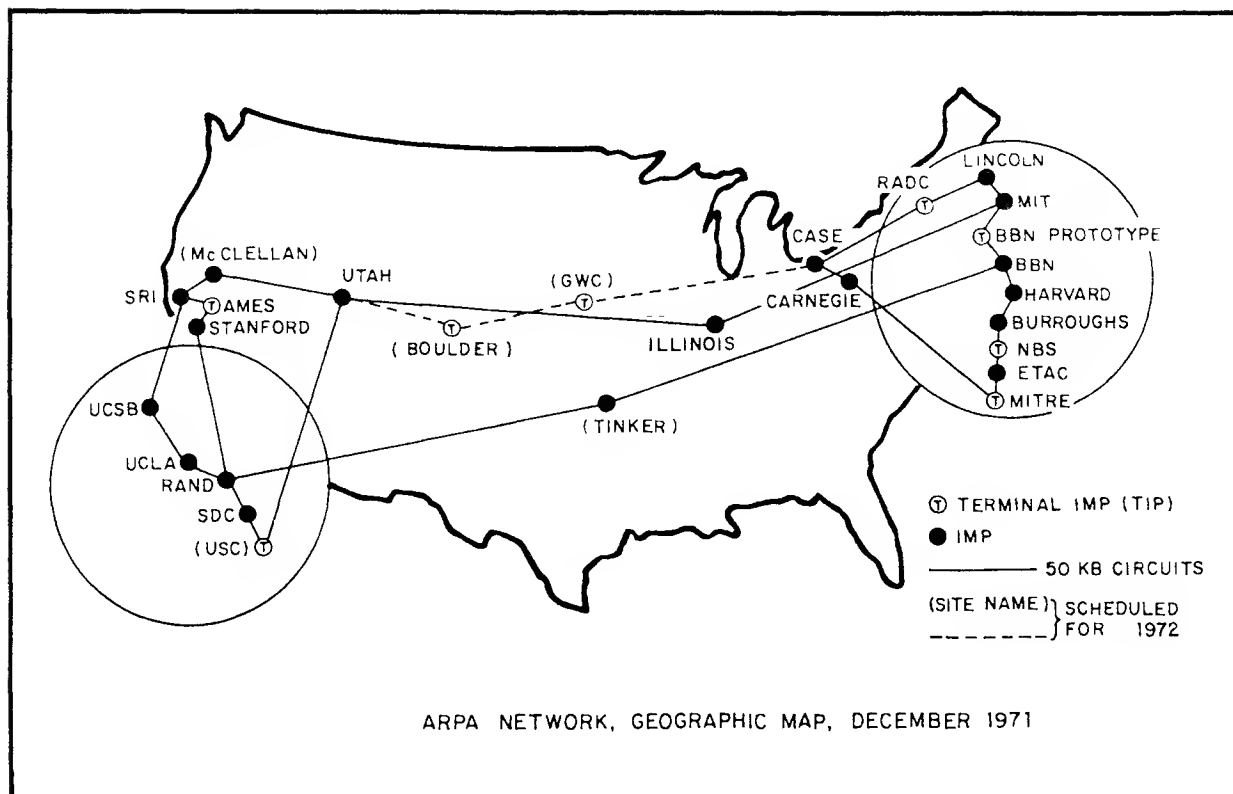
During the early 1960s, various research projects -- most of them funded by Arpa (the Advanced Research Projects Agency of the Defense Department) -- brought the phrase "timesharing" to the lips of computer industry buffs. This work paved the way for a business that will hit \$400 million this year, up 33% over 1971.

Now, Arpa is quietly touting another concept -- resource sharing -- that may lead to vast extensions of the early timesharing concept. By tying together a large number of computers through communications "nodes" -- themselves little computers -- the Arpa researchers have come up with a new kind of network . . . a community of computers that share vast calculation power, massive data stores, and specialized software.

The system, called Arpanet, already contains over 25 nodes (chart, page 3) and is looking for more users. Its final impact on the computer industry is unmeasurable at this time, but could be significant; the concepts behind Arpanet reflect many ideas in the mainstream of advanced planning for distributed computing and data transmission. Similar systems could be built for immediate commercial use.

In a nutshell, Arpanet links together computer resources so that they are used at maximum efficiency. Three aspects are involved:

- Load sharing brings a wealth of computer power to Arpa researchers on demand. Weather Bureau scientists can use the Sigma 7 at UCLA; MIT students, if they overload the school's PDP-10 during peak morning hours, can access a similar machine at Utah for the overflow, with the reverse possible after lunch to smooth the system load.
- Data base sharing prevents duplication of efforts. All researchers have access to the massive data bases generated at other sites, and this capability will increase next year when Illiac IV and its trillion-bit store join the network.
- Software sharing cuts one of the most significant costs in using a computer, and will provide transfer of human skills. Arpa researchers can access any specialized software on any system rather than develop their own or adapt it to suit a particular type of hardware.

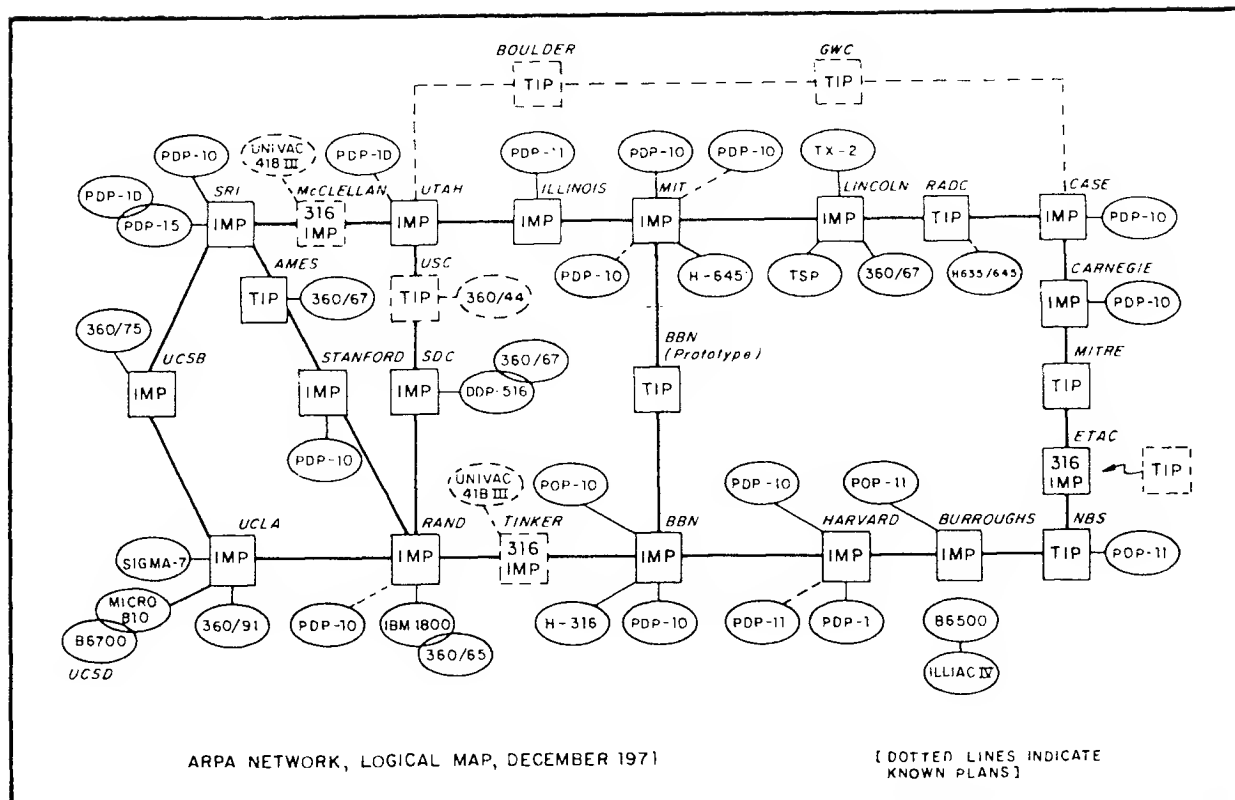


With such an interconnected system, virtually all the independent research efforts to advance computing in the U.S. are brought together under one roof, as it were, and will have a common orientation. No duplication of effort will be required. Yet all researchers interested will be able to access a one-of-a-kind machine such as Illiac IV, which joins the network next March, or work together on a common language for the system. This fusion of human efforts combined with almost human communication between machines is what Arpa means by resource sharing.

How Arpanet works

As the map above shows, Arpanet is a closed circuit network of 50 kilobit communications lines that connect over 25 research centers across the nation. Interconnection by conventional means would have required separate lines between each pair of nodes -- in a "star" fashion -- and the failure of any single line would have isolated those two nodes from each other. With just 40,000 miles of data transmission pipeline, however, Arpanet virtually guarantees communications capability between any two nodes at all times. If one link is unusable, another route always exists; it is located automatically by an element of the system.

The heart of the network is the node, and the network has many hearts. This little computer called an IMP (Interface Message Processor) receives and transmits messages according to a standard set of rules called a protocol. The IMP results from the system store-and-forward technology developed by Bolt Beranek and Newman Inc., also one of the timesharing originators. The node hardware is based on a Honeywell DP-516 or DDP-316 minicomputer that is modified, installed, and maintained by BBN. Each of the IMPs, priced at \$60,000, contains the system rules, or protocol, and they reload each other in the event of failure or removal from the communication network for maintenance.



The IMP is what actually connects the computers (Hosts) to the network. If a researcher at Carnegie, for example, wants to access a data base at Rand, his Host PDP-10 must transmit an appropriate message to its IMP. The IMP, in turn, acknowledges receipt of the message, translates it into a standard code, selects the fastest route -- considering the distance involved, the status of the lines, and load factors along the way -- and transmits the message to the next IMP in the route.

As the message proceeds to its destination, each intermediate IMP tells the previous one the message was received. Then, it checks the selected route to make sure it is still the best way to the final destination, and continues the transmission process. Ultimately the message arrives at its destination. In this example, for instance, the final IMP decodes the message and transmits it to the Rand Host. That Host then starts a confirmation message on its way back to the Carnegie Host to acknowledge receipt.

Amazingly, this entire process takes no more than 0.2 seconds. In fact, when the system was demonstrated to EDP Industry Report at the BBN control center recently, we saw no noticeable difference in transmission time from one Host to another. It didn't seem to matter whether we interacted with a Host in Utah or the Project MAC computer center down the street at MIT in Cambridge. The only apparent delays stemmed from complex operating systems in the Host computers themselves.

The coding and decoding tasks performed by the IMPs are no small feat. In fact, it is precisely for this reason that so many different types of computers can participate as Hosts. The schematic above gives some indication of the enormous complexity of the system; each different Host -- from Burroughs, DEC, IBM, XDS -- communicates in a different code. The IMP simply doesn't care.

Bring on the terminals

A more recent development within Arpanet extends the resource sharing concept still further and opens the available resources to multitudes of new users. Working as development coordinator for the system, Bolt Beranek and Newman has added to the IMP sufficient capability for interfacing a terminal directly to Arpanet. This modified IMP is called a TIP (for Terminal Interface Processor) and includes multiplexing capabilities for up to 63 terminals, costs \$120,000.

Now, for example, weather researchers at ETAC (Environmental Technical Applications Center) in Washington, D.C. can use computer power, software, and data at the University of Utah rather than having to duplicate these resources in Washington. Also, as the logical map (page 3) shows, Global Weather Central will soon be hooked into the system through another TIP.

The network, today, contains a relatively complete array of computer power, but is operating at only a few percent of capacity. Last fall, for example, the 675,000 message-per-day traffic was just over 2% of the 30 million capacity. So Arpa sees the next rash of new nodes as being TIPs rather than additional IMPs and associated Hosts.

How much does it cost?

The greatest single advantage in tying together computers in this fashion -- whether for research as in Arpanet or for certain business-related uses that are projected -- is the tremendous economy. Extensive studies undertaken for Arpa show that the capacity and cost of such distributed networks are "remarkably insensitive" to the distribution and destination of traffic. The total traffic is the only important parameter. Thus, says Arpa, "it is appropriate to charge for traffic initiated at a node, based on the cost of increasing the total capacity of the network by that amount." Charges, thus, could be much like those of the Postal Service, with message length (weight) more important than distance.

At present, an equal share of the communication line cost is being allocated to each node, so the cost for Arpanet users is about \$4,800 per node per month -- \$3,100 for communications and \$1,700 for leasing a minimal IMP. Use of a TIP increases the cost by \$1,600 a month.

But this will change. As soon as feasible, Arpa plans to charge for traffic initiated at a node. The actual charges will depend on the loading that can be expected for the system, since it will naturally not operate at peak capacity day in and day out. Arpa's target rate, based on 36% loading, is an estimated 30¢ per megabit in addition to equipment rental.

A few problems

As can be expected for any new development in computing, a number of problems -- legal, technical, and human -- will prevent this concept from mushrooming full bloom into a series of networks connecting most of the computers in the country. Five or ten years down the pike, however, significant changes in computer system organization are expected. And in the light of such a timetable, the anticipated problems don't seem insurmountable. In fact, they are quite similar to those already being faced by a number of computer users and, because of the resource sharing concept, if they are solved once, they are solved for all. Some of today's "challenges" are:

- Computer/communications interrelationships have been studied exhaustively by the Federal Communications Commission (EDP/IR Apr.9,1970), but all the questions weren't answered. It would appear the Arpanet would be what the FCC calls a "hybrid" system -- "an offering of service which combines Remote Access data processing and message-switching to form a single integrated service." Such systems, the FCC says, need be regulated if and only if message switching is the main purpose of the system and data processing is an incidental and integral part of the system. It looks to EDP/IR as if additional definition -- or at least a firmer ruling -- will be in the cards.
- Host-Host protocol is an on-going research project within the Arpanet members to allow graphics, block transfer of data, and file transfers. A common macro language doesn't seem to be necessary at this point in time, but it would appear necessary before similar systems gain widespread commercial usage.
- The privacy question will likely raise its head high as vast networks are established. The Arpanet really doesn't affect the invasion of privacy issue per se, but it could be a strong factor in bringing the question to the surface faster and forcing answers sooner than would be expected under current conditions.
- On the other hand, the sharing or resale of common carrier facilities cannot be done for profit. Thus, if several service organizations such as Control Data, GE, and Tymshare set up a jointly owned network (see below), provision for this current FCC requirement might be a problem.

The potential impact

As suggested earlier, many commercial "communities" today have characteristics and needs for which an Arpanet-like linkage might be useful today. There's no question but that some of these will evolve:

The timesharing industry itself could gain two primary advantages from setting up a jointly-owned system:

- + The network could provide a highly reliable communications system for interconnecting remote locations and central computers;
- + By mutual agreement, services within the network could draw on each other's powers and resources for backup and a more complete spectrum of service capabilities.

The large service companies mentioned above have their own private communications networks, but smaller companies cannot afford the luxury of such mammoth development efforts or the full time broadband communications links required. Also, even for these large companies, individual networks do not operate at nearly the efficiency that Arpanet sees right now.

Inter-airline reservations would be greatly simplified by interconnection through an Arpanet cousin. If one airline could not supply the requested seat, it could automatically query the network to find out who could do it, then write the ticket and perform the appropriate accounting transactions . . . all without any further human intervention.

Hospital patient records become increasingly dispersed as American society becomes more mobile. The Arpanet concept, however, provides a means for any participating hospital to query all other hospitals until a given patient's records are found. The local Host could then update the records on the remote Host rather than duplicating storage.

Inter-library lending systems could extend the resources of the nation's numerous libraries far beyond their present limits. Many companies have small, specialized libraries and draw on others for information they do not maintain themselves. At present, however, a librarian in one of these organizations must search manually for a library that has the publication required. A computerized network of library catalogs could track down the appropriate material far more quickly and efficiently.

Federal and state governments have numerous applications for which such a network could cut computer investment and line charges significantly. Some initial efforts at cooperation -- such as between motor vehicle bureaus -- are underway, but many other areas haven't even been considered.

Looking more than a few years into the future, EDP/IR can foresee even the possible specialization of computers by function. Certain computers, for example, might perform only payroll calculations for all network participants; another network might keep track of corporate shareholder records; and a third might automatically compute everyone's income tax and bill the individual for the amount due. Many of these would work with banks in the transfer of money. Arpa has expressed similar possibilities for the selection of computers:

Soon after a network with a dozen or more reliable computer services becomes available, many institutions will find it far more economic to obtain their computing services from a selected set of these remote systems, rather than run their own computer center.

For example, take the case of an institution about ready to upgrade its facility. One choice would be to obtain a medium scale, general purpose batch system, admittedly a compromise for their large numerical users and timesharing users, but the best single system that they can afford. Alternatively, they could buy no new machine and obtain access to several of the systems on the network through a TIP. This approach permits their large numerical users to use a large "number-cruncher," their statistical and payroll users to access a large scale general purpose system, and their interactive users to have Teletype or graphic console access to a good timesharing system. Overall the cost of each service is less than it would have been on a dedicated G-P computer by factors of between two and ten. Also, they can buy just the capacity they need and expand smoothly rather than having to pay for an oversize machine for a year or two.

That's hardware sharing, and the same type of advantages are expected from data base sharing and software sharing. Once these are going, says Arpa, there will develop other important network applications, "all of which require a large viable network before they become important in their own right. These include teleconferencing, publishing, library services, and office paperwork filing and distribution. Ten to twenty years from now these applications may well dominate computer usage and network usage but they are not likely to be important factors for at least five years." EDP/IR, along with Arpa, is convinced this is the wave of the future. And, in their petty pace, things are headed in this direction.

GENERAL-PURPOSE COMPUTERS CONCENTRATED AT LARGEST SITES IN THE UNITED STATES

With all the talk of "saturation" that goes around from time to time, EDP/IR asked the International Data Corporation, its publisher, to compile some relevant data. The figures below for general-purpose (Group A) computers only are based on the January, 1972 U.S. Computer Installation Data File maintained by IDC. Figures for tax reporting units -- plant sites for industrial companies -- are extrapolations based on 1968 data from 1970 Statistical Abstracts, U.S. Department of Commerce. This is the latest data available to IDC at presstime.

<u>No. Employees per tax reporting unit</u>	<u>Estimated No. of reporting units</u>	<u>Estimated % with 1 or more computers</u>	<u>No. of sites with one or more computers</u>	<u>Estimated total No. of Group A computers</u>
1- 19	3,260,000	0.05	1,630	1,630
20- 49	276,000	1.25	3,450	3,620
50- 99	95,400	4.7	4,500	4,950
100- 249	54,500	15	8,290	9,810
250- 499	17,550	46	8,000	10,215
500- 999	7,110	75	5,330	7,460
1000-2499	3,620	90	3,260	6,520
2500-4999	935	98	915	2,745
Over 5000	422	99.5	420	2,100
SUBTOTAL	3,713,537	0.96	35,795	49,050
Federal Government			1,325	3,585
State & Local Government			<u>2,310</u>	<u>2,725</u>
TOTALS			39,430	55,360

The figures above indicate that, on average, there are 1.4 general-purpose computers per site in the U.S. That average is slightly misleading, however, because 75% to 80% of all the sites have only one computer. The remainder, as a result, have almost 3 computers per site.

And this computer concentration is even more pronounced when the fact that some companies have many sites is considered. Because of this concentration some 40% of the computers in the U.S. are operated within only 700 or so organizations. Thus, as would be expected, computers follow wealth and high employment. The figures above indicate this concentration: Even though just under 1% of all tax reporting units have a computer, the penetration jumps to about 8% when the first category -- one-to-twenty employees -- is eliminated from the averages.

There are many opportunities for installing small computers at the number of small tax reporting units without EDP. But the way in which vast numbers of these organizations will receive computer power is probably going to be via terminals connected to a network of some sort. This is one more reason why commercial operations such as the one now up-and-running for Arpa researchers (page 1) will be so important. And other factors point in this direction. Shipments to new users accounted for about one-third of all computer deliveries during the mid-1960s; but these peaked in 1968 and 1969 -- at \$1.2 billion -- and have started to decline in value as well as percentage. By the mid-1970s, shipments to new users in the U.S. will be less than 10% of the total, according to IDC projections.

IBM'S NEW COMMUNICATIONS CONTROLLER/PROCESSOR HAS EXTENDED-TERM LEASE PLAN

In a move that further suggests its commitment to communications, IBM this week announced a controller "with its own processor for improved remote computing." The device -- called the 3705 -- is available for shipment this summer, and it is being offered "under terms of a new IBM extended term plan." Thus, in a single announcement treated routinely, IBM has eliminated a weak link in its product line and started people talking again about the possibility of IBM's offering term leases for additional products -- including mainframes.

All IBM will say about the new extended-term lease plan is that "we expect to announce the ETP to other products, but we will make that decision as the products come along. At present, we have no plans to add announced products to long term leasing plans."

The plan provides for discounts of about 13% off the regular month-to-month rental contract for the product. It provides for an initial 24-month contract and gives the user the option of an unlimited number of one-year extensions of the contract. And since communications systems are normally installed at considerable expense and planning, the two-year minimum would seem academic.

The new product will meet "growing requirements for more flexible and reliable teleprocessing operations." A "combination of logic circuitry and core memory," the 3705 is not dubbed a computer by IBM. But it will fill a gap that many IBM users, with the airlines at the van, have gone to minicomputer manufacturers to supply. IBM says it will take over the task of controlling the access of data to and from terminals in a network, thereby relieving the central computer of routine tasks such as communication line control, character checking, and buffering and line polling.

The 3705 can handle up to 352 communications lines, twice as many as the 2703 it replaces. It can handle line speeds ranging from 45.5 bps to 50,000 bps, and can support both asynchronous or binary synchronous transmission. The basic unit consists of a processor with 16K bytes of core memory. This can be expanded in 32K increments up to 240K bytes, or characters. The ETP charges range from \$1,200 to \$9,500 a month, compared to monthly charges of \$1,390 to \$10,940 under the shorter-term plan. Purchase is \$57,000 to \$449,000.

There is nothing very new, technically, about the 3705. Some observers have suggested that it is another home for the memory and logic circuits from 360s being returned by 370 customers. In any event, it broadens IBM's communications offerings and could cut into business previously received by Honeywell or Interdata from customers with 360-based communications systems.

AT&T SAYS DATA TRANSMISSION IS ITS FASTEST GROWING BUSINESS AT 50% A YEAR

The 1971 annual report just published by American Telephone & Telegraph Company contains a good deal of talk, and some figures, about the data communications business as seen by the giant. "Data transmission -- on private line channels and on our general message network -- continues, as it has for some years, to be our fastest growing service. It brought in revenues of some \$650 million in 1971 -- and market studies project a ten-fold increase in the decade ahead." AT&T's revenues from "data services" grew 18% last year, according to a table in the report, versus the 21% and 25% figures posted in the prior two years. The compound rate for the last ten years has been 26% annually.